

TRANSFER OF CDF DATA TO CNAF

Gene Oleynik & Dmitry Litvintsev
Fermilab

Assumptions

- 1 PB/mo (380 MB/s) can scale up or down, but this is an anchor
- Will be using resources migrating LTO4 to T10000C at the same time
- CDF Analysis on the dCache nodes will be active at the same time

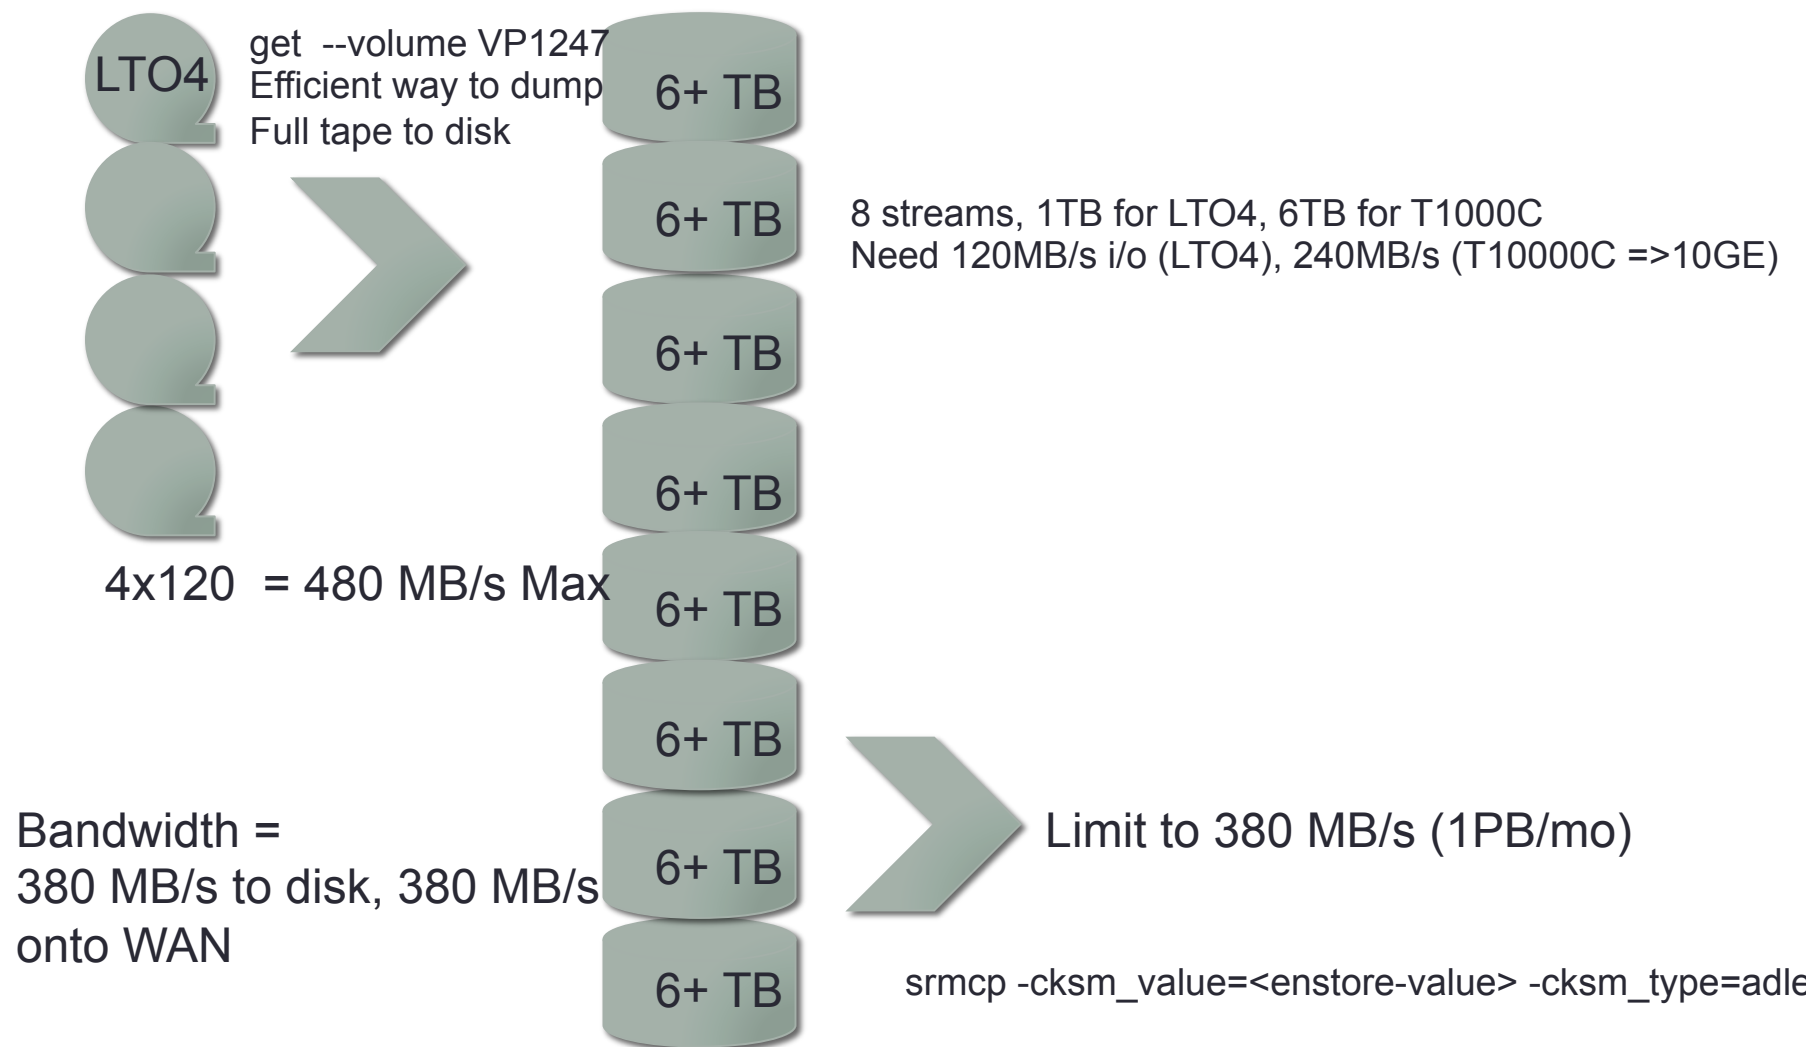
Two options to consider

- Push – Fermilab pushes the data to CNAF with dedicated machines
- Pull – CNAF sets up transfers from Fermilab dCache

Push

- Simple loop over tapes, dump to disk then write to network. Unit of work is tape, not file.
- Maximal use of resources: 8 streams, 4 dedicated LTO4 (2 T1000C), 4 moderate file servers (with 10GE for T10000C)
 - Can use efficient enstore get instead of encp to dump tape to disk
 - No disk spindle or tape drive contention so maximizes resources
- Can easily use srmcp with enstore checksum end-to-end
- How would CNAF sync this?
 - Can have hardware reflection at CNAF (e.g. 2 file servers serving say 2 T10000C)?

Push from Fermilab to CNAF



Pull by CNAF from CDF dCache

- Use existing Fermilab CDF dCache endpoint
- Controlled by CNAF
- A portion of the data is already cached (20%?) that could be moved more quickly,
- Could impact normal CDF dCache operations
 - 20% hit rate (already cached) x 380 MB/s = 76 MB/s impact on production dCache
 - Needs at least 8 nodes bandwidth (GE interfaces)
- Worst case need up to 11 dedicated LTO4 drives and may interfere with CDF migration and analysis.
 - Based on measured ave. LTO4 rate of 36 MB/s on CDF dCache
 - Taking into account already cached files may only need 8 drives
 - Pre-staging should improve but requires further investigation

Tests

- Recommend we explore both push and pull (Assuming CNAF can accommodate either)
 - Almost completed setting up server for push test
 - Can use CDF dCache pools being decommissioned for pull test?
- Dmitry has more slides on dCache option if time permits